connected to the cluster units,

the switching unit observes the MAC destination address of a packet arriving at said switching unit and checks if said MAC destination address is a MAC multicast address,

if so, said switching unit forwards the packet to those ports connected to the cluster units sharing a MAC multicast address.

REMARKS

The Examiner has objected to the drawings on grounds the processor 280 is missing from the drawings but mentioned in the specification. A proposed drawing correction is submitted herewith adding processor 280 to Figure 5.

The Examiner has also rejected claims 1-10 for lack of enablement under 35 USC 112, para. 1. The lack of enablement rejection is based upon the Examiner's assertion that no two devices can have the same MAC address, but the claimed invention recites the use of MAC multicast addresses.

It is the burden of the Examiner to make out a *prima facie* case of enablement, and that burden includes the need to show that undue experimentation would be required to make and use the claimed invention. In re Angstadt, 537 F.2d 489, 190 USPQ 214, 219 (CCPA 1976). The enablement requirement only requires that the specification be sufficiently detailed to enable one of ordinary skill in the art to make and use the invention without undue experimentation. Panduit Corp. v. Dennison Mfg. Co., 810 F.2d 1561, 1 USPQ2d 1593, 1595-6 (Fed. Cir. 1986). The specification need not disclose that which is well known in the art. Lindemann Maschinenfabrik GmbH v. American Hoise & Derrick Co., 730 F.2d 1452, 221 USPQ 481, 489 (Fed. Cir. 1984). Therefore, the elements of the *prima facie* case of enablement are:

1. A rational basis as to

    a. why the disclosure does not teach or

    b. why to doubt the objective truth of the statements in the disclosure that purport to teach

2. the manner and process of making and using the invention

3. that corresponds in scope to the claimed invention

4. to one of ordinary skill in the pertinent technology,

5. without undue experimentation, and

6. dealing with subject matter that would not already be known to the skilled person as of the filing date of the application.

The Examiner must provide evidence from the specification support each of these elements for an enablement rejection to be proper.

To combat the *prima facie* case of enablement, the applicants must prove that at least one element of the case has not been met. Applicant's argument is directed to element 6. MAC multicasting and MAC multicast addresses are well known terms and concepts known by persons of ordinary skill in the art. As evidence of this fact, consider Exhibit A attached hereto. This is a printout from a web site at the University of Aberdeen and is a short description of the MAC protocol. This description shows at page 2 that MAC addresses in Ethernet^{tm} local area networks can be multicast addresses. Although this article of Exhibit A is dated 1/1/2001, Exhibit B proves that Ethernet and MAC multicast addresses have been in the prior art since the early 70's when Ethernet was invented at Xerox PARC. Exhibit B is an excerpt from the treatise Comer, *Internetworking With TCP/IP: Principles, Protocols, and Architecture*, pp. 20 and 29 (Prentice Hall). Page 2 of Exhibit B proves that Ethernet has been around since the

early 70's and page 3 of Exhibit B proves that Ethernet addresses are also known as MAC addresses and that MAC addresses can be multicast addresses.

Thus, the Examiner should now be convinced that the calls in the claims for MAC multicast addresses are well understood by those skilled in the art and that such multicast addresses can and do exist for use by the invention.

The Examiner also rejected claims 1-10 for indefiniteness on grounds the phrase in claims 1-6 because the phrase "specific to the cluster" is not understood. In response to this rejection, the phrase "specific to the cluster" has been removed from claims 1 and 6.

Favorable action is ernestly solicited.

Respectfully submitted,

Dated: July 15, 2004

Ronald Craig Fish
Reg. No. 28,843
Tel 408 778 3624
FAX 408 776 0426

I hereby certify that this correspondence is being deposited with the United States Postal Service as First Class Mail, postage prepaid, in an envelope addressed to: Commissioner for Patents , P.O. Box 1450, Alexandria, Va. 22313-1450.

on ____________________
(Date of Deposit)

Ronald Craig Fish, President
Ronald Craig Fish, a Law Corporation
Reg. No. 28,843

UNIVERSITY OF ABERDEEN Department of Engineering

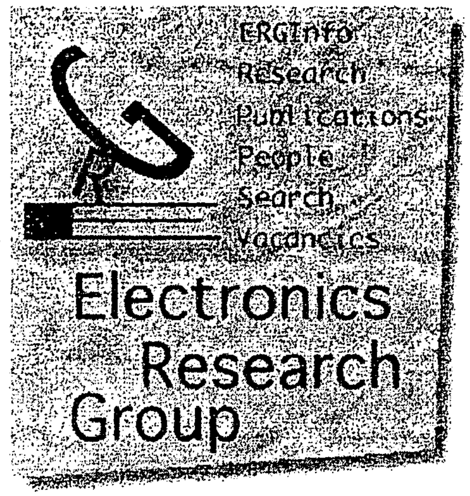

[ERG info] [Research] [Publications]
[People] [Search] [Vacancies]

Comments on this Web Site: webmaster@erg.abdn.ac.uk

You appear to be visiting this site from the **IPv4** address :- **194.100.2.65**

Powered by APACHE

EXHIBIT A -1-
5 PAGES

Medium Access Control (MAC)

The Medium Access Control (MAC) protocol is used to provide the data link layer of the Ethernet LAN system. The MAC protocol encapsulates a SDU (payload data) by adding a 14 byte header (Protocol Control Information (PCI)) before the data and appending a 4-byte (32-bit) Cyclic Redundancy Check (CRC) after the data. The entire frame is preceded by a small idle period (the minimum inter-frame gap, 9.6 microsecond (µS)) and a 8 byte preamble.

Preamble

The purpose of the idle time before transmission starts is to allow a small time interval for the receiver electronics in each of the nodes to settle after completion of the previous frame. A node starts transmission by sending an 8 byte (64 bit) preamble sequence. This consists of 62 alternating 1's and 0's followed by the pattern 11. Strictly speaking the last byte which finished with the '11' is known as the "Start of Frame Delimiter". When encoded using Manchester encoding, at 10 Mbps, the 62 alternating bits produce a 5 MHz square wave.

The purpose of the preamble is to allow time for the receiver in each node to achieve lock of the receiver Digital Phase Lock Loop which is used to synchronise the receive data clock to the transmit data clock. At the point when the first bit of the preamble is received, each receiver may be in an arbitrary state (i.e. have an arbitrary phase for its local clock). During the course of the preamble it learns the correct phase, but in so doing it may miss (or gain) a number of bits. A special pattern (11), is therefore used to mark the last two bits of the preamble. When this is received, the Ethernet receive interface starts collecting the bits into bytes for processing by the MAC layer.

Header

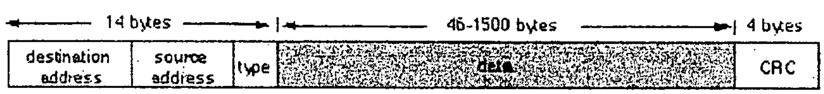

MAC encapsulation of a packet of data

The header consists of three parts:

- A 6-byte destination address, which specifies either a single recipient node (unicast mode), a group of recipient nodes (multicast mode), or the set of all recipient nodes (broadcast mode).
- A 6-byte source address, which is set to the sender's globally unique node address. This may be used by the network layer protocol to identify the sender, but usually other mechanisms are used (e.g. arp). Its main function is to allow address learning which may be used to configure the filter tables in a bridge.
- A 2-byte type field, which provides a Service Access Point (SAP) to identify the type of protocol being carried (e.g. the values 0x0800 is used to identify the IP network protocol, other values are used to indicate other network layer protocols). In the case of IEEE 802.3 LLC, this may also be used to indicate the length of the data part.

CRC

The 32-bit CRC added at the end of the frame provides error detection in the case where line errors (or transmission collisions in Ethernet) result in corruption of the MAC frame. Any frame with an invalid CRC is discarded by the MAC receiver without further processing. The MAC protocol does not provide any indication that a frame has been discarded due to an invalid CRC.

Inter Frame Gap

EXH A -2-

After transmission of each frame, a transmitter must wait for a period of 9.6 microseconds (at 10 Mbps) to allow the signal to propagate through the receiver electronics at the destination. This period of time is known as the Inter-Frame Gap (IFG). While every transmitter must wait for this time between sending frames, receivers do not necessarily see a "silent" period of 9.6 microseconds. The way in which repeaters operate is such that they may reduce the IFG between the frames which they regenerate.

Byte Order

It is important to realise that nearly all serial communications systems transmit the least significant bit of each byte first at the physical layer. Ethernet supports broadcast, unicast, and multicast addresses. The appearance of a multicast address on the cable (in this case an IP multicast address, with group set to the bit pattern 0xxx xxxx xxxx xxxx xxxx xxxx) is therefore as shown below (bits transmitted from left to right):

```
0                             23 IP Multicast Address Group  47
|                             | <--------------------------->|
1000 0000 0000 0000 0111 1010 xxxx xxx0 xxxx xxxx xxxx xxxx
|                                      |
Multicast Bit                          0 = Internet Multicast
                                       1 = Assigned for other uses
```

However, when the same frame is stored in the memory of a computer, the bits are ordered such that the least significant bit of each byte is stored in the right most position (bits transmitted right-to-left within octets, octets transmitted left-to-right):

```
0                             23                             47
|                             |                              |
0000 0001 0000 0000 0101 1110 0xxx xxxx xxxx xxxx xxxx xxxx
        |                     <--------------------------->
     Multicast Bit              IP Multicast Address Group
```

CSMA /CD

The Carrier Sense Multiple Access (CSMA) with Collision Detection (CD) protocol is used to control access to the shared Ethernet medium.

Receiver Processing Algorithm

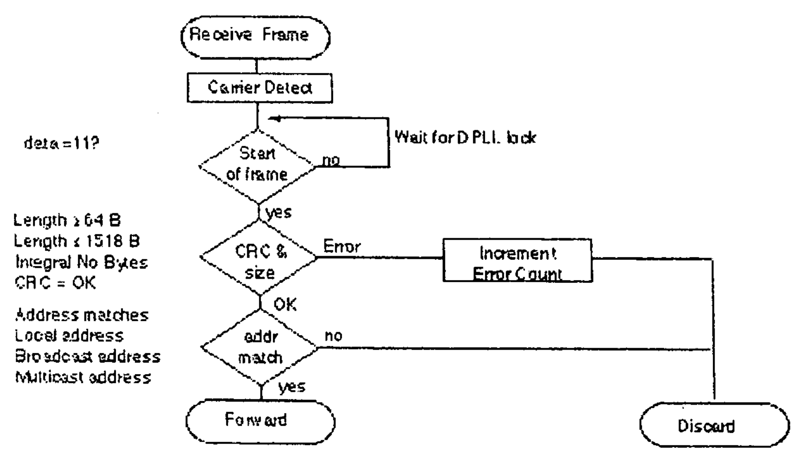

A Runt Frame

Ex4. A -3-

Any frame which is received and which is less than 64 bytes is illegal, and is called a "runt". In most cases, such frames arise from a collision, and while they indicate an illegal reception, they may be observed on correctly functioning networks. A receiver must discard all runt frames.

A Giant Frame

Any frame which is received and which is greater than the maximum frame size, and is called a "giant". In theory, the jabber control circuit in the transceiver should prevent any node from generating such a frame, but certain failures in the physical layer may also give rise to over-sized Ethernet frames. Like runts, giants are discarded by the Ethernet receiver.

A Misaligned Frame

Any frame which does not contain an integral number of received octets (bytes) is also illegal. A receiver has no way of knowing which bits are legal, and how to compute the CRC-32 of the frame. Such frames are therefore also discarded by the Ethernet receiver.

Other Issues

The Ethernet standard dictates a minimum size of frame, which requires at least 46 bytes of data to be present in every MAC frame. If the network layer wishes to send less than 46 bytes of data the MAC protocol adds sufficient number of zero bytes (0x00, is also known as null padding characters) to satisfy this requirement. The maximum size of data which may be carried in a MAC frame using Ethernet is 1500 bytes (this is known as the MTU in IP).

A protocol known as the "Address Resolution Protocol" (arp) is used to identify the MAC source address of remote computers when IP is used over an Ethernet LAN.

Exception to the Rule

An extension to Ethernet, known as IEEE 802.1p allows for frames to carry a tag. The tag value adds an extra level of PCI to the Ethernet frame header. This increases the size of the total MAC frame when the tag is used. A side effect of this is that NICs and network devices designed to support this extension require a modification to the jabber detection circuit.

See also

- Calculations involving Ethernet
- CSMA/CD and Ethernet Transmit Algorithm
- Ethernet
- Ethernet Packet Decodes
- IP packet processing
- MAC Address Assignment
- Manchester Coding
- Network Interface Cards

ExH A -4-

Gorry Fairhurst - email:G.Fairhurst@eng.abdn.ac.uk - Date: 01/01/2001 EG3557

EXH A -5-

Internetworking With TCP/IP
Vol I:
Principles, Protocols, and Architecture
Fourth Edition

DOUGLAS E. COMER

Department of Computer Sciences
Purdue University
West Lafayette, IN 47907

PRENTICE HALL
Upper Saddle River, New Jersey 07458

EXH. B 3 PAGES

[redacted]

2.3.1 Network Hardware Addresses

Each network hardware technology defines an *addressing mechanism* that computers use to specify the destination for a packet. Every computer attached to a network is assigned a unique address, usually an integer. A packet sent across a network includes a *destination address field* that contains the address of the intended recipient. The destination address appears in the same location in all packets, making it possible for the network hardware to examine the destination address easily. A sender must know the address of the intended recipient, and must place the recipient's address in the destination address field of a packet before transmitting the packet.

Each hardware technology specifies how computers are assigned addresses. The hardware specifies, for example, the number of bits in the address as well as the location of the destination address field in a packet. Although some technologies use compatible addressing schemes, many do not. This chapter contains a few examples of hardware addressing schemes; later chapters explain how TCP/IP accommodates diverse hardware addressing schemes.

2.4 Ethernet Technology

Ethernet is the name given to a popular packet-switched LAN technology invented at Xerox PARC in the early 1970s. Xerox Corporation, Intel Corporation, and Digital Equipment Corporation standardized Ethernet in 1978; IEEE released a compatible version of the standard using the standard number *802.3*. Ethernet has become the most popular LAN technology; it now appears in virtually all corporate networks as well as many small installations. Because Ethernet is so popular, many variants exist. Although the original wiring scheme has been superceded, understanding the original design helps clarify the intent and some of the design decisions. Thus, we will discuss the original design first, and then cover variants.

Formally known as *10Base5*, the original Ethernet design uses a coaxial cable as Figure 2.1 illustrates.

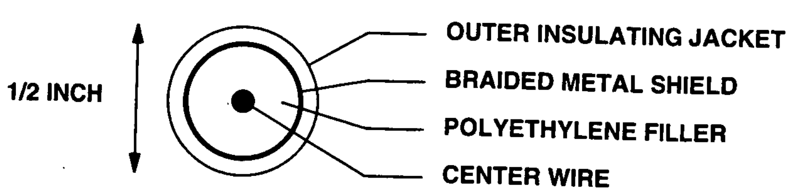

Figure 2.1 A cross-section of the coaxial cable used in the original Ethernet.

Called the *ether*, the cable itself is completely passive; all the active electronic components needed to make the network function are associated with the computers attached to the network. Each Ethernet cable is about 1/2 inch in diameter and up to 500

EXH. B -2-

2.4.9 Ethernet Hardware Addresses

Ethernet defines a 48-bit addressing scheme. Each computer attached to an Ethernet network is assigned a unique 48-bit number known as its *Ethernet address*. To assign an address, Ethernet hardware manufacturers purchase blocks of Ethernet addresses† and assign them in sequence as they manufacture Ethernet interface hardware. Thus, no two hardware interfaces have the same Ethernet address.

Usually, the Ethernet address is fixed in machine readable form on the host interface hardware. Because each Ethernet address belongs to a hardware device, they are sometimes called *hardware addresses*, *physical addresses*, *media access (MAC) addresses*, or *layer 2 addresses*. Note the following important property of Ethernet physical addresses:

> *Physical addresses are associated with the Ethernet interface hardware; moving the hardware interface to a new machine or replacing a hardware interface that has failed changes the machine's physical address.*

Knowing that Ethernet physical addresses can change will make it clear why higher levels of the network software are designed to accommodate such changes.

The host interface hardware examines packets and determines the packets that should be sent to the host. Recall that each interface receives a copy of every packet that passes through a hub — even those addressed to other machines. The host interface uses the destination address field in a packet as a filter. The interface ignores those packets that are addressed to other machines, and passes to the host only those packets addressed to it. The addressing mechanism and hardware filter are needed to prevent a computer from being overwhelmed with incoming data. Although the computer's central processor could perform the check, doing so in the host interface keeps traffic on the Ethernet from slowing down processing on all computers.

A 48-bit Ethernet address can do more than specify a single destination computer. An address can be one of three types:

- The physical address of one network interface (a *unicast address*)
- The network *broadcast address*
- A *multicast address*

By convention, the broadcast address (all 1s) is reserved for sending to all stations simultaneously. Multicast addresses provide a limited form of broadcast in which a subset of the computers on a network agree to listen to a given multicast address. The set of participating computers is called a *multicast group*. To join a multicast group, a computer must instruct its host interface to accept the group's multicast address. The advantage of multicasting lies in the ability to limit broadcasts: every computer in a multicast group can be reached with a single packet transmission, but computers that choose not to participate in a particular multicast group do not receive packets sent to the group.

†The Institute for Electrical and Electronic Engineers (IEEE) manages the Ethernet address space and assigns addresses as needed.

EXH. B -3-